Effective use of the PGAS Paradigm: Driving Transformations and Self-Adaptive Behavior in DASH-Applications

Kamran Idrees
High Performance Computing Center Stuttgart (HLRS)
idrees@hlrs.de

Tobias Fuchs
Ludwig-Maximilians-Universität München (LMU)
tobias.fuchs@nm.ifi.lmu.de

Colin W. Glass
High Performance Computing Center Stuttgart (HLRS)
glass@hlrs.de

arXiv:1603.01536v1 [cs.DC] 4 Mar 2016

Abstract

DASH is a library of distributed data structures and algorithms designed for running the applications on modern HPC architectures, composed of hierarchical network interconnections and stratified memory. DASH implements a PGAS (partitioned global address space) model in the form of C++ templates, built on top of DART – a run-time system with an abstracted tier above existing one-sided communication libraries.

In order to facilitate the application development process for exploiting the hierarchical organization of HPC machines, DART allows to reorder the placement of the computational units. In this paper we present an automatic, hierarchical units mapping technique (using a similar approach to the Hilbert curve transformation) to reorder the placement of DART units on the Cray XC40 machine Hazel Hen at HLRS. To evaluate the performance of new units mapping which takes into the account the topology of allocated compute nodes, we perform latency benchmark for a 3D stencil code. The technique of units mapping is generic and can be be adopted in other DART communication substrates and on other hardware platforms.

Furthermore, high–level features of DASH are presented, enabling more complex automatic transformations and optimizations in the future.

Categories and Subject Descriptors D.1.2 [*PROGRAMMING TECHNIQUES*]: Automatic Programming; D.1.3 [*PROGRAMMING TECHNIQUES*]: Concurrent Programming

Keywords DASH, DART, PGAS, UPC, MPI, Self-Adaptation, Code Transformations

1. Introduction

Partitioned Global Address Space (PGAS) devises a method of parallel programming by introducing a unified global view of address space (like purely shared memory systems) and presiding over the distribution of data (similar to a distributed memory system), in order to provide a programmer with ease of use and a locality-aware paradigm. To distribute the data across the system, PGAS implementations use one–sided communication substrates, which are hidden from the application developer.

Unified Parallel C (UPC) is an implementation of the PGAS model. UPC has a single shared address space, which is partitioned among UPC threads, such that a portion of shared address space resides in the memory local to a UPC thread. UPC provides mechanisms to distinguish between local and remote data accesses, thus allowing to capitalize on data locality. However, a programmer needs to perform custom coding, potentially even building advanced data distribution schemes, to exploit locality efficiently.

DASH is a C++ library that delivers distributed data structures and algorithms designed for modern HPC machines, which are well-suited for hierarchical network interconnections and memory stratum (Fürlinger et al. 2014). DASH aims at various domains of scientific applications, providing the programmer with advanced data structures and algorithms, consequently reducing the need for custom coding.

DASH calls DASH Run-Time (DART), which provides basic functionalities to build a PGAS model using state of the art one-sided communication libraries (Zhou et al. 2014). These functionalities include:

- Global memory management and optimization for accessing data that reside on shared memory system (Zhou et al. 2015)
- Creation, destruction and management of teams and groups
- Collective and non-collective communication routines
- Synchronization primitives

This paper presents DASH as an alternative to traditional PGAS implementations like UPC. Short-comings of UPC and other traditional PGAS implementations are discussed, which in many cases prevent an effective use of the PGAS paradigm. Furthermore, features of DASH overcoming these short–comings are presented. The main contributions of this paper are:

1. We present an advanced local copy feature in DASH

2. We evaluate the throughput of the local copy feature
3. We present an automatic hierarchical units mapping mechanism for applications having a nearest neighbor communication pattern
4. We evaluate the applicability of automatic hierarchical units mapping mechanism on a 3D stencil communication kernel on Cray XC40 Hazel Hen machine at HLRS

2. Experiences with UPC

From our experience with UPC for our in–house molecular dynamics code, we highlighted three major issues, resulting in severe performance degradation (Idrees et al. 2013). These issues – and a further problem regarding hardware topology – are:

1. Manual pointer optimization is necessary for fast access to local data (using local pointer)
2. Non-trivial data distribution schemes need to be implemented by hand
3. Communication is performed at the same granularity as data access
4. No mechanism available for co–locating strongly interacting units on the given hierarchical hardware topology

The first issue regards the failure of UPC compilers to automatically distinguish between shared and distributed memory data accesses, even though the complete data layout is available. This holds true for both static and dynamic allocation of data (as the block size of a distributed shared array needs to be a compile–time constant). This can lead to a significant performance degradation and can only be avoided by expert programmer intervention. Manual optimization requires checking all parts of the code where significant data accesses are performed and switching to local pointers for local memory access.

The second issue is that UPC provides only round robin and blocked data distribution schemes. These schemes are suboptimal for many applications featuring some sort of short range geometric data accessing patterns, e.g. stencil patterns. This will lead to a unnecessarily high amount of communication traffic and percentage of remote communication. To avoid this problem, the programmer has to write specific data mapping routines.

The third issue is associated with the communication granularity. In shared memory address space, the programmer can directly access and modify the shared data. The PGAS paradigm also provides these attributes for its global address space. However, accessing and modifying remote data is expensive, especially if it leads to many small communications. As UPC does not change the granularity, this often leads to a vast number of tiny communications. To avoid this problem, the programmer needs to take the underlying distributed memory architecture into account and perform the necessary optimizations for packing communications manually.

The fourth issue addresses the difficulty of adapting the behavior of an application to the machine topology. For example, reordering the placement of software units that may allow to reduce the communication cost by placing the interacting partners closer to each other on the physical hardware.

To summarize: in order to achieve a near optimal performance using traditional PGAS implementations, the programmer needs to take care of a variety of issues manually. This contradicts the driving idea behind PGAS: ease of programmability. The good news is, DASH is tackling these issues by providing automatic optimization for faster local data accesses (Zhou et al. 2015), advanced data distribution schemes (Fuchs et al. 2015), algorithm specific routines for pre-fetching and packing of data and automatic hierarchical units mapping. The following section provides a detailed illustration of these advance features.

3. DASH as a Solution

DASH resolves the short–comings of the traditional PGAS implementations with its automatic optimizations, advanced data structures and algorithms. We will now explain a few specific features of DASH which address the problems highlighted in the previous section.

3.1 Fast Access to Local Data

The automatic detection of local vs. remote data access in DASH demonstrates an effective use of PGAS paradigm: every data access is performed in the most efficient way available. This automatic behavior is achieved by capitalizing on the shared memory window feature of MPI-3 (Hoefler et al. 2012) used in the MPI version of DART (DART–MPI). The shared memory window can be accessed directly by local MPI processes using load/store operations (zero–copy model), allowing the processes to circumvent the single-copy model of the MPI layer. DART-MPI maps both global and shared memory windows to the same shared memory region, thus allowing the DART units on shared memory to directly access the local memory region. The DART units that are not part of the shared memory window, perform RMA operations using a global window. Furthermore, the use of the zero copy model for intra-node communication in DART scales down the memory bandwidth problem. We have demonstrated in (Zhou et al. 2015), our optimization of mapping both shared and global memory windows to the same memory region on shared memory, enables faster intra-node communication. This allows DASH programmers to iterate over distributed data structures without worrying about slow local data accesses, unlike UPC where manual pointer optimizations are necessary to avoid less efficient local data accesses (Idrees et al. 2013).

3.2 High-Level Data Distribution Schemes

DASH features several data distribution schemes (*patterns*) that provide highly flexible configurations depending on data extents and logical unit topology. New pattern types are continuously added to DASH. This flexibility leads to a large number of data distributions that can be used for a single use case.

The preferable pattern configurations depend on the specific use case. Algorithms operating on global address space strictly depend on domain decomposition as they expect data distributions that satisfy specific properties.

Without methods that help to configure data distributions, programmers must learn about the differences between all pattern implementations and their restrictions. We therefore provide high-level functions to automatically optimize data distribution for a given algorithm and vice-versa. These mechanism are described in detail in (Fuchs et al. 2015). In this we present a classification of data distribution schemes based on well-defined properties in the three mapping stages of domain decomposition: partitioning, mapping, and memory layout. This classification system serves two purposes:

- Provides a vocabulary to formally describe general data distributions by their semantic properties (*pattern traits*).
- Specifies constraints on expected data distribution semantics.

As an example, the *balanced partitioning* property describes that data is partitioned into blocks of identical size. An algorithm that is optimized for containers that are evenly distributed among units can declare the *balanced partitioning* and *balanced mapping* properties as constraints. A mapping is balanced if the same number of blocks is mapped to every unit.

When applying the algorithm on a container, its distribution is then checked against the algorithm's constraints already at compile time to preventing inefficient usage:

```
static_assert(
   dash::pattern_contraints<
       pattern,
       partitioning_properties<
         partitioning_tag::balanced >,
       mapping_properties<
         mapping_tag::balanced >,
       layout_properties< ... >
   >::satisfied::value);
```

Listing 1: Checking distribution constraints at compile time and run time.

Finally, pattern traits also allow to implement high-level functions that resolve data distribution automatically. For this, we use simple constrained optimization based on type traits in C++11 to create an instance of a initially unspecified pattern type that is optimized for a set of property constraints. To be more specific, the automatic resolution of a data distribution involves two steps: at compile time, the pattern type is deduced from constraints that are declared as type traits. Then, an optimal instantiation of this pattern type is resolved from distribution constraints and run-time parameters such as data extents and team size.

```
TeamSpec<2> logical_topo(16, 8);
SizeSpec<2> data_extents(extent_x, extent_y);
// Deduce pattern:
auto pattern =
  dash::make_pattern<
    partitioning_properties<
      partitioning_tag::balanced >,
    mapping_properties<
      mapping_tag::balanced >
  >(sizespec, teamspec);
```

Listing 2: Deduction of an Optimal Distribution using dash::make_pattern.

Deduction of data distribution can also interact with team specification to find a suitable logical Cartesian arrangement of units. As for domain decomposition, DASH provides traits to specify preferences for logical team topology such as "compact" or "node-balanced". As a result, application developers only need to state a use case, such as DGEMM, and let unit arrangement and data distribution be resolved automatically.

While no automation can possibly do away with the need for manual optimization in general, automatic deduction as provided by DASH greatly simplifies finding a configuration that is suitable as a starting point for performance tuning. In comparison, finding practicable blocking factors and process grid extents for ScaLAPACK, even in seemingly trivial use cases, is a challenging task for non-experts.

3.3 Creating Local Copies

Algorithms and container types in DASH follow the concepts and semantics of their counterparts in the C++ Standard Template Library (STL) and are consequently based on the iterator concept. Algorithms provided by the STL can also be applied to DASH containers and most have been ported to DASH providing identical semantics. Programmers will therefore already be familiar with most of the API concepts. For copying data ranges, the standard library provides the function interface `std::copy`. In DASH, a corresponding interface `dash::copy` is provided for copying data in PGAS.

This section presents the concept of the the functions `dash::copy` and `dash::copy_async`, first-class citizens in the DASH algorithm collection which represent a uniform, general interface for copy operations within global address space.

As an example, compare how an array segment is copied using standard library and using DASH:

```
double range_copy[ncopy];
// STL variant:
std::copy(arr.begin(), arr.begin() + ncopy,
          range_copy);
// DASH variant, on a global array:
dash::copy(arr.begin(), arr.begin() + ncopy,
           range_copy);
```

Asynchronous variants of data movement operations are essential to enable overlap of communication and computation. They employ the *future* concept also known from the standard library:

```
double range_copy[ncopy];
auto fut_copy_end = dash::copy_async(
                      arr.begin(),
                      arr.begin() + ncopy,
                      range_copy);
// Blocks until completion and returns result:
double * copy_end = fut_copy_end.wait();
```

Copying data from global into to local memory space is a frequent operation in PGAS applications. It can involve complex communications pattern as some segments of the copied range might be placed in memory local to the requesting unit's core while others are owned by units on distant processing nodes.

Performance of copy operations in partitioned global address space is optimized by avoiding unnecessary data movement and scheduling communication such that interconnect capacity is optimally exploited. We use the following techniques in DASH, among others:

Shared memory For segments of the copied data range that are located on the same processing node as the destination range, `std::copy` is used to copy data in shared memory. This reduces calls to the communication backend to the unavoidable minimum. This is not restricted to copying: DASH algorithms in general automatically distinguish between accesses in shared and distributed memory. And even when not using DASH algorithms, the DASH runtime automatically resorts to shared window queries and `memcpy` instead of MPI communication primitives for data movement within a processing node.

Chunks To achieve optimal throughput between processing nodes, communication of data ranges is optimized for transmission buffer sizes by splitting the data movement into chunks. Adequate chunk sizes are obtained from auto tuning or interconnect buffer sizes provided in environment variables by some MPI runtimes.

Communication scheduling Parallel transmission capacity is exploited whenever possible: if the data source range spans the address space of multiple units, separate asynchronous transmissions for every unit are initiated instead of a sequence of blocking transmissions. Also, the single asynchronous operations are then ordered in a schedule such that communication is balanced among all participating units to fully utilize interconnect capacity.

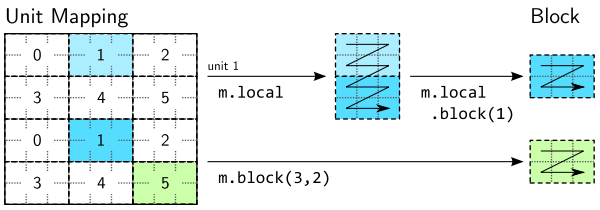

Figure 1: Accessing blocks in a two-dimensional DASH matrix. The `local` access modifier changes the scope of the succeeding block access to local memory space.

The shared memory optimization technique require means to logically partition a global data range into local and remote segments. For this, DASH provides the utility function `dash::local_range` that partitions a global iterator range into local and remote sub-ranges:

```
dash::Array<T> a(SIZE);
// Get iterator ranges of local segments in
// global array:
std::vector<dash::Range<T>> l_ranges =
  dash::local_range(a.begin() + 42, a.end());
// Iterates local segments in array:
for (l_range : l_ranges) {
  // Iterates values in local segments:
  for (l_value : l_range) {
    // ...
  }
}
```

The global iterator range to partitioned into locality segments may be multidimensional. In addition, DASH containers provide methods to access blocks mapped to units directly so that programmers do not have to resolve partitions from domain decomposition themselves. For example, sub-matrix blocks can be copied in the following way:

```
dash::Matrix<2, double> matrix;
// First block in matrix assigned to
// active unit, i.e. in local memory space:
auto l_block = matrix.local.block(0);
// Specify matrix block by global block
// coordinates, i.e. in global memory space:
auto m_block = matrix.block({ 3, 5 });
// Copy local block to global block:
dash::copy(l_block.begin(),
           l_block.end(),
           m_block.begin());
```

The `dash::copy` function interface and the `dash::Matrix` concept greatly simplify the implementation of linear algebra operations. Efficiency of the underlying communication is achieved without additional effort of the programmer due to the optimization techniques presented in this section.

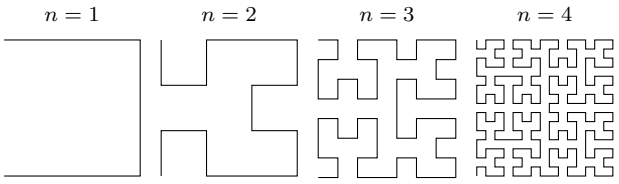

Figure 2: Hilbert Space Filling Curve Example

3.4 Automatic Hierarchical Units Mapping

The mapping of an application's software units (or threads/processes/tasks) to the physical cores on an HPC machine is becoming increasingly important due to the rapid increase in the number of cores on a machine (Hoefler et al. 2011)(Deveci et al. 2014). This also leads to increasingly hierarchical networks and – depending on the underlying system and the submitted job – sparse core allocations. Therefore, if two units of an application which are interacting partners (or communicating more frequently than average) are placed far from each other in the network, they will have to communicate through several levels in the network hierarchy. The placement of these units not only has repercussions in their communication latency and bandwidth, but may also result in the congestion of the network links.

As the units mapping plays a vital role, DART-MPI provides DASH with the mean to automatically reorder the units which respect both the communication pattern of an application and the topology of allocated nodes on an HPC machine. This results in reduced communication and overall execution time of an application. Currently we counter this problem for a specific set of applications which are based on nearest neighbor communication. The programmer only needs to inform DASH that the application to be executed has a nearest neighbor communication pattern. The automatic mapping routine then gathers the required hardware topology, computes a new hierarchical unit mapping and registers it in the system. The new mapping is determined based on an approach similar to a Hilbert Space Filling Curve (HSFC) partitioning (Moon et al. 2001). An example of HSFC is shown in figure 2. HSFC is chosen due to its property of preserving the locality. The algorithm however does not fix the length (number of elements to iterate over in a dimension) of HSFC for multiple levels as the lengths are dependant upon the number of nodes corresponding to each level in the network hierarchy.

Before discussing the steps in the automatic hierarchical unit mapping algorithm, we briefly explain the hierarchical network levels of the Cray XC40 Supercomputer Hazel Hen in the following.

3.4.1 Off-Node Network Hierarchy on Hazel Hen

The first off-node network hierarchical level is a *compute blade*. A compute blade is composed of four nodes which share an Aries chip. The Aries chip connects these four nodes with the network interconnect. This is the fastest connection between nodes.

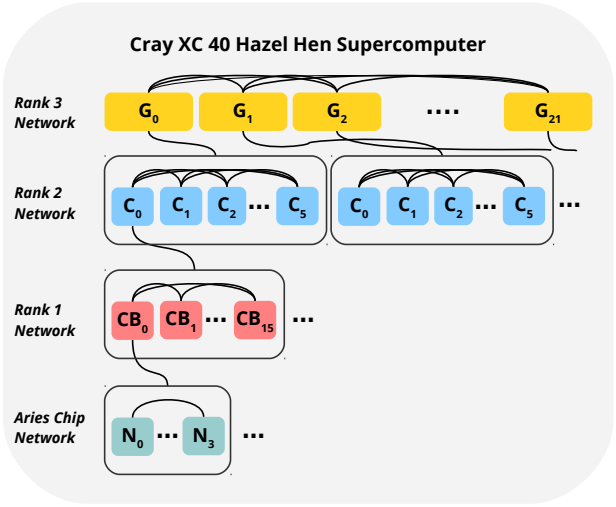

Figure 3: Layered Layout of the Network Hierarchy of `Hazel Hen`. G, C, CB and N are abbreviated for Group, Chassis, Compute Blade and Node respectively.

The second level is the *rank 1 network* (or backplane). The rank 1 network is used for inter-compute blades communication within a chassis – set of 16 compute blades (64 nodes). The rank 1 network has adaptive routing per packet. Any two nodes at this level communicate with each other by going first through their Aries chip, then through the backplane, and finally to the Aries chip of the target node. This is an all-to-all PC board network.

The *rank 2 network* is used for inter-backplane communication (nodes on distinct chassis) in a two cabinet group (a group is composed of 384 nodes). The backplanes are connected through copper cables. All copper and backplane signals run at 14 Gbps. The minimal route between two nodes on distinct chassis is two hops, whereas the longest route requires four hops. The aries adaptive routing algorithm is used to select the best route from four routes in a routing table. The rank 2 network also has an all-to-all connection, connecting 6 Chassis (two cabinets).

The last off-node network level is the *rank 3 network*, which is used for communication between different groups. The rank 3 network has all-to-all routing using optical cables. If minimal path between two groups is congested, traffic can be hopped through any other intermediate group (1 or 2 hops).

The layered layout of network hierarchy of Hazel Hen is show in figure 3.

3.4.2 Algorithm

The automatic hierarchical units mapping algorithm is executed by every DASH unit and assumes that the user has performed a binding of the units to the CPU cores, such that the units do not migrate from one CPU to another [1]. Every unit performs the following steps:

1. Acquires the *total number of units* in the team (to be mapped on the hierarchical topology).
2. Acquires its processor name and parses it to obtain the *Node ID* on which the unit resides.
3. Participates in the collective allgather operation to obtain the Node IDs of all units (units on the same node have the same Node ID) and uses the Node ID as a key to look for the placement information string of the node inside a topology file of the machine [2].
4. Reads topology file to acquire placement information string, number of sockets and number of cores per socket, for each allocated node.
5. Parses the placement information string of each node in order to obtain the value of each hierarchical level of the machine corresponding to each node.
6. Sorts the nodes with respect to all levels in the *network hierarchy* i.e. at first performing the sorting according to the values of every node on Level[4], then on Level[3] and so on. For example:

   Level(4, 3, 2, 1, 0) = (0, 0, 1, 12, 3)
   Level(4, 3, 2, 1, 0) = (0, 0, 1, 13, 0)
   Level(4, 3, 2, 1, 0) = (0, 0, 1, 13, 1)
   ⋮
   Level(4, 3, 2, 1, 0) = (9, 1, 2, 1, 1)
7. Determines balanced distribution of total number of units in a cartesian grid.
8. Performs balanced distribution of units per node, to form multi–core groups in order to reduce inter-node communication (For example: a balanced distribution for 24 cores as in Hazel Hen would be $4 \times 3 \times 2$. The lengths of coordinate directions of the 3D Cartesian grid (x,y,z) of total number of units should be divisible by the cartesian grid of units per node (e.g. (4,3,2)). This is necessary as our reordering method (Algorithm 1) performs multi–core grouping at the node level and therefore the number of groups in each coordinate direction should fit the cartesian grid of total number of units.

[1] On a Cray machine, unit or process binding can be easily performed by adding the argument *-cc cpu* to the *aprun* command.

[2] A topology file on a Cray machine can be created using Cray's *xtprocadmin* utility. The placement information string of a node looks like *c11-2c0s15n3*, which means the position of the node in the machine hierarchy is: column 11, row 2, chassis 0, compute blade 15 and node 3.

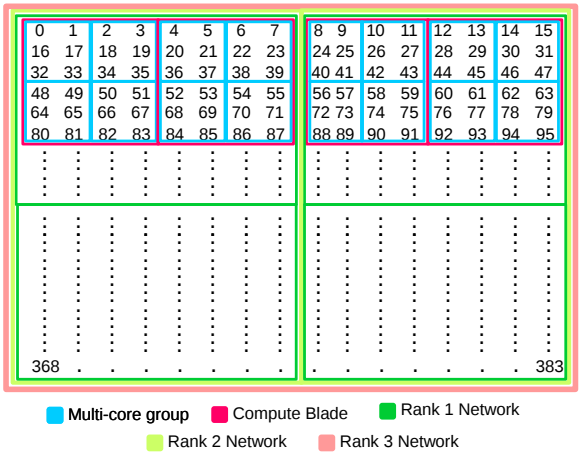

Figure 4: Example of automatic hierarchical units mapping for a 2D nearest neighbor problem of size (24×16). Please note that the example is just for elaborating the hierarchical units mapping and does not reflect the actual number of hardware instances at each network hierarchy level.

9. Assigns new unit ID to each unit taking into consideration the multi–level network hierarchy, i.e. multicore groups of units are mapped as close as possible in the network hierarchy in order to reduce communication between distinct network hierarchy levels.
10. Finally, the reordered unit IDs are registered in the system.

After the last step, the new mapping is completed. The algorithm will result in an optimal units mapping if the node allocation is contiguous on all network hierarchy levels. Optimal being the minimal surface area, which results in minimal communication traffic on all network hierarchy levels. If the nodes are allocated in a sparse manner, the algorithm attempts to preserve the locality through its HSFC–like implementation. Figure 4 shows an example of automatic hierarchical units mapping for a 2D nearest neighbor communication pattern.

4. Performance Evaluation of Local Copy

Optimization techniques employed in copying data in global memory space have been discussed in 3.3. In the following, we evaluate the *local copy* use case where a unit creates a local copy of a data range in global memory. We consider the following scenarios, named by the location of the copied data range:

local Both source- and destination range are located at the same unit. This scenario does not involve communication as `dash::copy` resorts to copying data directly in shared memory. To illustrate the maximum achievable throughput, this scenario is also evaluated using `std::copy`.

Algorithm 1: Pseudo–code of algorithm to compute new unit IDs which respect underlying machine's network hierarchy

Input : Balanced distribution of total number of units, number of units per node

Output: Unit IDs respecting network hierarchy

```
1  unitNumber ← 0
2  N4 ← numThirdLevelBlocksInFourthLevel[3]
3  N3 ← numSecondLevelBlocksInThirdLevel[3]
4  N2 ← numFirstLevelBlocksInSecondLevel[3]
5  N1 ← numNodesAtFirstLevel[2]
6  N0 ← numUnitsPerNode[3]
7  NT ← numTotalUnits[3]
   /* Given that the nodes are sorted according
      to multiple network hierarchy levels,
      compute the new unit IDs                    */
8  foreach x, y, z: 0 → N4 (0, 1, 2) do
9      foreach a, b, c: 0 → N3 (0, 1, 2) do
10         foreach d, e, f: 0 → N2 (0, 1, 2) do
11             foreach g, h: 0 → N1 (0, 1, 2) do
12                 foreach i, j, k: 0 → N0 (0, 1, 2) do
13                     if unitNumber > totalUnits then
14                         break
15                     end
16                     fourthLevelOffset ← ..
17                     thirdLevelOffset ← ..
18                     secondLevelOffset ← d × N0[0] ×
                       NT[1] × NT[2] + e × N1[0] × N0[1]
                       × NT[2] + f × N1[1] × N0[2];
19                     firstLevelOffset ← g × N0[1] × NT[2]
                       + h × N0[2];
20                     unitID[unitNumber] ←
                       fourthLeveOffset + thirdLevelOffset +
                       secondLevelOffset + firstLevelOffset +
                       i × NT[1] × NT[2] + j × NT[2] + k;
21                     unitNumber++;
22                 end
23             end
24         end
25     end
26 end
   /* Register the new unit IDs in run time
      system -- reordering                        */
```

socket The data range to be copied and the copy target range are owned by units mapped to different sockets on the same processing node. In this case, communication is avoided by DART recognizing the data movement as an operation on shared memory windows.

remote The source range is located on a remote processing unit and is copied in chunks using `MPI_Get`.

For meaningful measurements, it is essential to avoid a pitfall regarding cache effects: if the copied data range has been initialized by a unit placed on the same processing node as the unit creating the local copy, the data to be copied is stored in L3 data cache shared by both units. In this case, the `local` and `socket` scenarios would effectively measure cache bandwidth instead of the more common and less convenient case where copied data is not available from cache.

4.1 Benchmark Environment

The local copy benchmark has been executed on SuperMUC phase 2 nodes for the available MPI implementations Intel MPI, IBM MPI, and OpenMPI. The MPI variants each exhibit specific advantages and disadvantages:

The installation of IBM MPI does not support MPI shared windows, effectively disabling the optimization in the DASH runtime for the *socket* scenario, but offers the most efficient non-blocking RDMA.

Intel MPI requires additional polling processes for asynchronous RDMA which increases overall communication latency.

The benchmark application has been compiled using the Intel Compiler (icc) version 15.0. Apart from being linked with different MPI libraries, the build environment is identical for every scenario.

4.2 Results

The results from all scenarios for the three MPI implementations is shown in Figure 5.

As a first observation, performance of `std::copy` varies with the MPI implementation used. This is because different C standard libraries must be linked for the respective MPI library. This also explains why cache effects become apparent for different range sizes in the *local* scenarios. In general, performance of local copying is expected to decrease for ranges greater than 32 KB which is the capacity of L1 data cache on SuperMUC Haswell nodes. The C standard library linked for OpenMPI sustains better performance for larger data sizes compared to the other evaluated MPI variants.

When copying very small ranges in the *local* scenario the constant overhead in `dash::copy` introduced by index calculations outweighs communication cost. Still, the employed shared memory optimization leads to improved throughput compared to MPI operations used in the *remote* scenario.

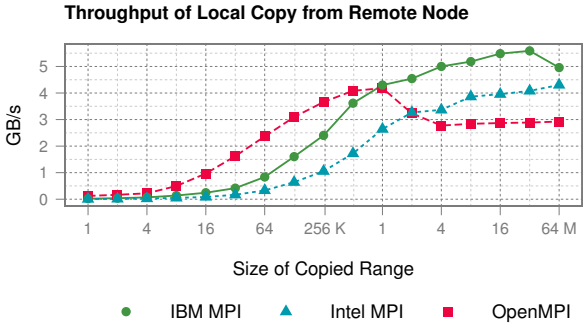

Figure 6: Throughput of `dash::copy` for local copies of blocks on a remote processing node.

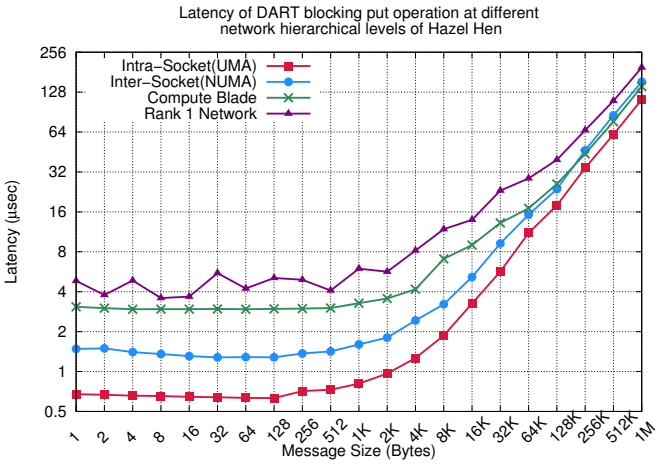

Figure 7: Average latency of transferring a message between two nodes at different hierarchical levels of the network. The message size is varied from 1 Byte to 1 MegaByte.

For copied data sizes of roughly 64 KB and greater, `dash::copy` achieves the maximum throughput measured using `std::copy`. This corresponds approximately to a minimum of a 90×90 block of double-precision floating point values and is far below common data extents in real-world use cases.

As expected, achieved throughput in the *socket* and *remote* scenarios are comparable for IBM MPI as shared window optimizations are not available and MPI communication is used to move data within the same node. Fortunately, IBM MPI also exhibits the best performance for MPI Communication. For ranges of 1 MB and larger, there is no significant difference between local and remote copying.

It might seem surprising that throughput in the *socket* scenario, where data is copied between NUMA domains, exceeds throughput in scenario *local* in some cases. However, in the *socket* scenario, data is copied in the DASH runtime using `memcpy` instead of `std::copy`. The different low-level variants are expected to yield different performance and again depend on the C standard library linked.

Figure 6 summarizes achieved throughput in the *remote* scenario of all MPI implementations in a single plot for comparison.

Results from this micro-benchmark can serve to autotune partition sizes used to distribute container elements among units. For example, a minimum block size of 1 MB is preferable for IBM MPI while block sizes between 1 to 16 MB should be avoided for Intel MPI and OpenMPI as NUMA effects decrease performance otherwise.

5. Performance Evaluation of the Stencil Kernel

We now evaluate the performance of 3D stencil communication kernel with and without using the automatic hierarchical units mapping feature. In order to measure solely the impact of hierarchical units mapping on the performance, we have disabled the shared memory window feature of DART-MPI for the benchmark shown later.

5.1 Evaluation Metric

In this stencil communication kernel, each unit communicates with six neighbors (left, right, upper, lower, front, and back). We use the blocking DART put operation for transferring the messages from one unit to another and the size of the messages is varied exponentially from 1 byte to 2 megabytes.

We are interested in evaluating the *relative performance improvement factor*, which is computed by taking the ratio of the average execution times (ten–thousand iterations) of stencil communication kernel using default (as performed by the job launcher on the Cray machine) against hierarchical units mapping.

5.2 Benchmark Environment

The benchmarks are carried out on the Cray XC40 machine Hazel Hen at HLRS. Each node on Hazel Hen is based on Intel Xeon CPU E5-2680 v3 (30M Cache, 2.50 GHz) processors and comprises 24 cores (12 cores per socket). Cray's Aries interconnect provides node-node connectivity with multiple hierarchical network levels. We have two on node memory hierarchy levels, which are Uniform Memory Access (UMA) – intra-socket communication – and Non-uniform Memory Access (NUMA) – inter-socket communication. It's easy to exploit the on node hierarchical levels. In this paper we are more interested in showing results by exploiting the off-node network hierarchical levels (Section 3.4.1).

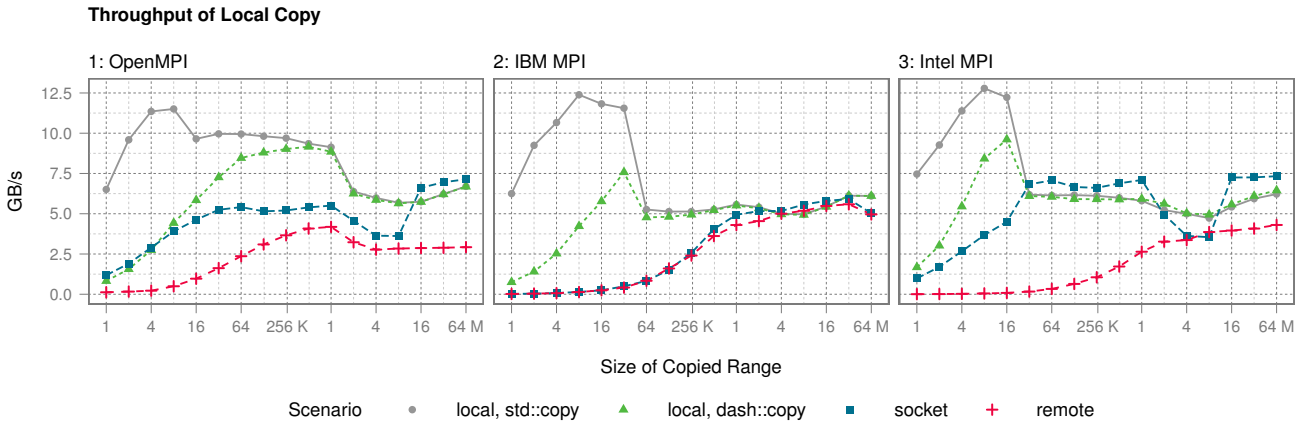

Figure 5: Throughput of `dash::copy` for local copies of data ranges from different locality domains. Throughput of `std::copy` is included for reference.

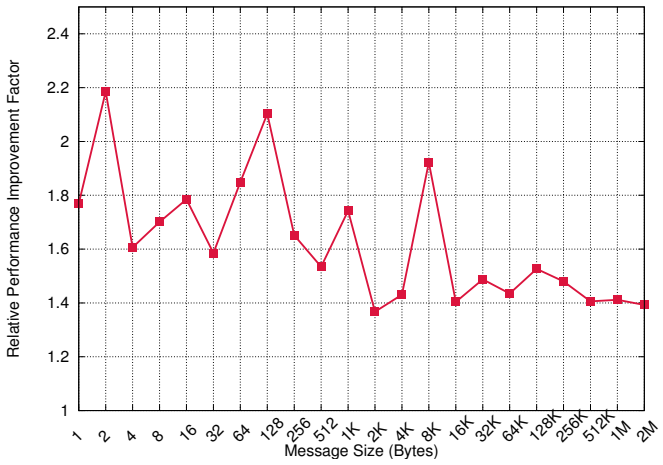

Figure 8: Relative performance improvement factor of 3D stencil communication kernel using default against automatic hierarchical units mapping, on 384 sparse nodes on Cray Hazel Hen XC40 machine. The message size is varied from 1 Byte to 2 MegaBytes.

5.3 Results

Figure 7 shows the average latency of messages up to the rank 1 network of Hazel Hen, highlighting the impact of different network levels.

Figure 8 shows the relative performance improvement factor of the 3D stencil communication kernel on 384 nodes (9,216 units) of Hazel Hen. The nodes were allocated in a sparse manner by Cray's job launcher, having small contiguous blocks of nodes. It can be seen that our units mapping provides an average performance improvement by a factor of 1.4 to 2.2.

6. Conclusions

In this paper we have presented specific features of DASH which resolve some issues we observed in traditional PGAS implementations. We have shown in section 5 that our automatic hierarchical units mapping provides a notable performance improvement over default units mapping. A user can take advantage of this self–adapting behavior without putting any effort into understanding the complex machine hierarchy or performing any custom coding. Furthermore, new features are presented enabling a user to represent the computation and communication patterns of scientific applications at a very high level of abstraction in DASH, while DASH will take care of necessary code transformations. We are currently working on extending methods for automatic data distribution to data flow scenarios. However, automatic optimization in many data flow use cases is conceptually equivalent to integer programming and thus proven to be NP-hard. We assume that solutions for a useful subset of scenarios can be found using linear programming techniques like the Simplex algorithm.

Acknowledgments

This work was supported by the project *DASH* which is funded by the German Research Foundation (DFG) under the priority program "Software for Exascale Computing - SPPEXA" (2013-2015).

References

Karl Fürlinger, Colin Glass, Jose Gracia, Andreas Knüpfer, Jie Tao, Denis Hünich, Kamran Idrees, Matthias Maiterth, Yousri Mhedheb, and Huan Zhou. Dash: Data structures and algorithms with support for hierarchical locality. In *Euro-Par 2014: Parallel Processing Workshops*, pages 542–552. Springer, 2014.

Kamran Idrees, Christoph Niethammer, Aniello Esposito, and Colin W Glass. Performance evaluation of unified parallel c for molecular dynamics. In *Proceedings of the 7th International Conference on PGAS Programming Models*, page 237.

Huan Zhou, Kamran Idrees, and José Gracia. Leveraging MPI-3 Shared-Memory Extensions for Efficient PGAS Runtime Systems In *Euro-Par 2015: Parallel Processing*, pages 373–384. Springer, 2015.

Huan Zhou, Yousri Mhedheb, Kamran Idrees, Colin W Glass, José Gracia, and Karl Fürlinger. DART-MPI: An MPI-based Implementation of a PGAS Runtime System In *Proceedings of the 8th International Conference on Partitioned Global Address Space Programming Models*, page 3. ACM, 2014.

Torsten Hoefler, James Dinan, Darius Buntinas, Pavan Balaji, Brian W Barrett, Ron Brightwell, William Gropp, Vivek Kale, and Rajeev Thakur. *Leveraging MPIs one-sided communication interface for shared-memory programming*. Springer, 2012.

Tobias Fuchs and Karl Fürlinger. Expressing and Exploiting Multidimensional Locality in DASH. *Springer Lecture Notes in Computational Science and Engineering*. Springer, November 2015.

Bongki Moon, Hosagrahar V Jagadish, Christos Faloutsos, and Joel H Saltz. Analysis of the clustering properties of the hilbert space-filling curve. *Knowledge and Data Engineering, IEEE Transactions on*, 13(1):124–141, 2001.

Torsten Hoefler and Marc Snir. Generic topology mapping strategies for large-scale parallel architectures. In *Proceedings of the international conference on Supercomputing*, pages 75–84. ACM, 2011.

Mehmet Deveci, Sivasankaran Rajamanickam, Vitus J Leung, Kevin Pedretti, Stephen L Olivier, David P Bunde, Umit V Catalyurek, and Karen Devine. Exploiting geometric partitioning in task mapping for parallel computers. In *Parallel and Distributed Processing Symposium, 2014 IEEE 28th International*, pages 27–36. IEEE, 2014.